

QUASI-ANONYMOUS CHANNELS

Ira S. Moskowitz

Center for High Assurance Computer Systems - Code 5540
Naval Research Laboratory, Washington, DC 20375, USA

&

Richard E. Newman

CISE Department

University of Florida, Gainesville, FL 32611-6120, USA

&

Paul F. Syverson

Center for High Assurance Computer Systems - Code 5540
Naval Research Laboratory, Washington, DC 20375, USA

Abstract

Although both *anonymity* and *covert channels* are part of the larger topic of *information hiding*, there also exists an intrinsic linkage between anonymity and covert channels. This linkage was illustrated in [1]; however, [1] just scratched the surface of the interplay between covert channels and anonymity, without a formal analysis of the related issues. This paper begins the process of formalizing the linkage between anonymity and covert channels via the study of *quasi-anonymous channels*. We also discuss and contrast some of the existing formal mathematical models of anonymity.

Key Words

Anonymity, mix, covert channel, information theory.

1. Introduction

Information hiding, amongst other things, contains two topics of interest, anonymity and covert channels. At first glance, these two topics seem to have very little to do with each other, but upon deeper analysis, we see that there is a linkage between them. This linkage has two aspects: one is that an “anonymity” system may be utilized to leak information “out of the system” via a covert channel; the other aspect is that a mathematical measurement of the covert channel as a communication channel can be used to measure the degree of anonymity that the anonymity system provides.

We note that network traffic analysis has also been shown to be exploitable via covert channels, e.g. [2]. While [2] did not specifically deal with the scenario of anonymous network traffic, in [3] an attempt was made to quantify prevention of network traffic covert channels by making the communication anonymous, thus extending the notions of [2] to the realm of anonymity.

This paper reviews anonymity and the various threat models, thus illustrating the pragmatic difficulties of complete anonymity. Therefore, in many cases the most we can hope for is quasi-anonymity (a fact well known by the anonymity community; although the words ‘anonymity’ and ‘anonymous’ are commonly used). It is this quasi-anonymity that allows covert communication channels with non-trivial ‘throughput’. We call such a covert channel a *quasi-anonymous channel*. How to measure quasi-anonymity is a matter of contention. We feel that a study of the quasi-anonymous channel addresses this problem. We propose to measure information leakage in an anonymity system by using the characteristics of the associated quasi-anonymous channel.

A covert channel is simply a communication channel that exists, contrary to system design, in a computer system or network [4], [5]. Covert channels have been well studied. The standard metric for a covert channel is its *capacity* [6]. However, capacity alone does not suffice for all covert channels (e.g., the *small message criterion* [5]), a lesson that may be applicable to our study of quasi-anonymous channels. However, in this paper we mainly concern ourselves with the optimal asymptotic error-free throughput of a covert channel, which is the capacity.

The standard anonymity terminology has been given by [7]: “*Anonymity* is the state of being not identifiable within a set of subjects, the *anonymity set*.” When a sender is anonymous within a set of potential senders, [7] defines “sender anonymity.” They go on to state that, “*Unobservability* is the state of IOIs [items of interest] being indistinguishable from any IOI at all.” Further, [7] follows with “This means that messages are not discernible from ‘random noise’.” They also state that “sender unobservability \Rightarrow sender anonymity.” In this paper, we explore the idea that random noise provides perfect anonymity. We also provide concrete examples

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2003		2. REPORT TYPE		3. DATES COVERED 00-00-2003 to 00-00-2003	
4. TITLE AND SUBTITLE Quasi-Anonymous Channels				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory,Center for High Assurance Computer Systems,4555 Overlook Avenue, SW,Washington,DC,20375				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 6	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

that may appear at first glance to be “anonymous,” but upon further analysis are not.

Our brief paper serves, we hope, two purposes. One is to quantify the lack of anonymity in a supposedly anonymizing network via a Shannon-type analysis. The other purpose is to show how a malicious user in an anonymity network may leak information, in a covert manner, to an eavesdropper on the network who is in cahoots with the malicious user. We realize that the first purpose is novel, but the second is a well-founded concern [4], [5].

2. Anonymous Networks---Classical Mix

A traditional way for obtaining anonymity is by the use of a mix [8]. The idea of a mix is to obfuscate which sender is sending to which receiver. Actual message content is hidden via various cryptographic means and is not a concern. What is a concern is the ability to de-couple sender from receiver. The desire is to keep this association hidden from an eavesdropper Eve (a “bad gal” adversary who is attempting to figure out who is transmitting to whom). If Eve has knowledge of who is sending and who is receiving without an active attack, she is a global passive adversary (GPA). But, if she only has partial knowledge of message traffic on the network, she is a restricted passive adversary (RPA).

A single mix is a single point of failure; therefore, an anonymous network may employ a chain of mixes. There are various ways that a mix may forward a message on, known as the *flushing algorithm* [9] of the mix. The “classical” mix is a threshold mix, where after a given number of messages have entered the mix, the mix then fires and flushes out all of its messages. This way Eve may have seen a message go in, but she (hopefully) cannot link the incoming message with an outgoing message. A timed mix may have a clock, and after a given allotment of time, the mix will flush all of its messages onto their destinations. Variants of these might involve the concept of a pool mix. In a pool mix, a certain number of messages are left behind in a random manner. But while pooling assists in confusing Eve, pooling has the disadvantage of increased message latency. The examples that on which we concentrate have a mix acting as an exit or entry point firewall to a private enclave and are extensions of classical mix theory.

Mix Firewalls --- Dual Enclave Scenario

In [1] we discussed using mixes as firewalls. The mixes in [1] are actually timed mixes. We take two enclaves. Communication within an enclave is private. Senders in Enclave1 wish to send messages to Enclave2. We assume that Eve can only count the encrypted messages between the enclaves and know the size of Enclave1.




Figure 1. Dual Enclave

We will often use this scenario of mixes as both outgoing firewalls and incoming firewalls (and the associated assumptions); we refer to this as *dual enclave*. We feel that this is an extremely important type of private communication between private enclaves. We must make sure we have a full analysis of what can be leaked out by quasi-anonymity. An example of this type of communication was given in [1]: packets from one LAN/enclave are sent to another LAN/enclave using IPSEC tunneling. Here, an eavesdropper Eve can only count the number of outgoing messages destined for the receiving LAN/enclave. In general, any situation where two private LANs wish to communicate over a public line applies. We will vary the type mixes that are used. In our following initial example, we use a timed mix with flushing time of 1 t (time unit).

Initial Example: There is a set time unit t , called a tick. Every t a sender in Enclave1 either sends or does not send a message to a receiver in Enclave2. Different senders may send to the same receiver. A sender can only send to one receiver. Before leaving Enclave1, the messages go through a mix firewall. Every t the mix fires and flushes all of its messages “in the open” to the second mix firewall. The messages are encrypted and appear identical to Eve. As assumed, the only thing Eve can do is to count messages. The second mix then forwards (flush time is not germane) the messages to the proper receivers in Enclave2. How anonymous is this scenario? Again, our concern is mostly with sender anonymity. Can Eve tell who sent a message? (We caution the reader that this dual enclave configuration is different than what is normally considered.) We are keeping hidden who is sending a message. Often anonymity analysis allows Eve to know the senders for a given mix batch, at least up to a probabilistic level. If there is only one sender in Enclave1, then there is no anonymity since Eve knows this fact. Eve simply sees when this lone sender is transmitting by counting 0 or 1 messages. What if there are two senders in Enclave1? Eve still has some knowledge of who is sending. Again Eve counts the messages across the public line. If there are no messages, Eve knows that no one is transmitting. If there are two messages, then Eve knows that both senders are transmitting. If there is one message, then Eve is confused. If Eve has some probabilistic knowledge of the behavior of senders in Enclave1, then she can do better than total confusion if she counts one message. If the senders in Enclave1 have the same probabilistic behavior, then Eve is in the state of maximal confusion, but

anonymity is still compromised. This is why we use the term ‘quasi-anonymity’. Of course, as the number of transmitters in Enclave1 increases, so does the quasi-anonymity of this network. A covert channel can utilize this quasi-anonymity to leak (from Enclave1) information to Eve. Similarly, the covert channel can be used to quantify the quasi-anonymity in this scenario. Study of this was initiated in [1].

3. Covert Channel Analysis

As discussed above, a covert channel is a communication channel that exists, contrary to system design, in a computer system or network. The covert channels that arise due to less-than-perfect anonymity are called *quasi-anonymous channels*. The channels are analyzed using Shannon’s information theory [6]. The simplest case of such a covert channel is a discrete memoryless channel (DMC). With respect to the above dual enclave example we assume that there is a malicious sender Alice in Enclave1. Alice wishes to covertly communicate with Eve by affecting Eve’s message count.

The lack of perfect anonymity is what enables Alice to communicate non-trivially with Eve.

We measure this anonymity by the amount of information that Alice can send to Eve. The maximum error-free amount of information per unit time that can be sent is given by the channel *capacity*. Note that the capacity may give the worst-case analysis for leakage, but one may wish to study sub-optimal amounts of information flow. This may be more properly studied by the mutual information, or perhaps by even simpler characteristics of the covert channel from Alice to Eve.

All that Alice can do is send or not send one message (it does not matter to which receiver Alice is transmitting--- Eve gets the same count). Therefore, Alice is represented by the random variable A , 0 is the input symbol corresponding to Alice not sending a message, and 0^c corresponds to Alice sending a message. We have the distribution $P(A=0) = x$, $P(A=0^c) = 1-x$. We must now assign a distribution to the other senders C_i in Enclave1. In [1] we assume that the other “clueless” senders C_i , $i = 1$ to N , are described by an i.i.d. with $P(C_i \text{ sends a message}) = q$. We assume that Alice and the C_i act independently of each other. Eve is described by the distribution E , the message count. The symbols that Eve may receive via the covert channel are $\{0,1,...,N+1\}$, depending on whether Alice is transmitting, and whether C_i is transmitting. Certainly, no matter what the probabilities, if Eve receives $N+1$, Eve knows that Alice is transmitting. Similarly, if Eve receives 0, Eve knows that Alice is not transmitting. For the values $\{1,...,N\}$ Eve can only perform a probabilistic guess. How can one call this anonymous communication? In some cases Eve knows exactly what Alice has done, and in other cases she can make a probabilistic guess. Even if Alice and the C_i all have the same probabilistic behavior, we cannot say that we have

anonymity. Thus, random noise does not guarantee anonymity.

We denote the mutual information between Alice and Eve as $I(A,E) = H(E) - H(E|A)$, the difference between the entropy of E , and the entropy of E conditioned on A , (see [1], [6]). The capacity C is the maximum of $I(E,A)$ as x varies from 0 to 1, with q fixed. The situation that minimizes the capacity is when C_i behaves as a random fair coin toss: $q = 1/2$. In fact, the value of x that maximizes $I(E,A)$ when $q = 1/2$, is when $x = 1/2$ also. This capacity is not zero [6]. As N approaches infinity the capacity goes to zero, but it is never zero. However, at first glance with every sender in Enclave1 having a 50-50 chance of sending a message, it might seem that we would have perfect anonymity.

We include in detail the special case of Alice and C (only one C_i) in Enclave1. The channel matrix describing the various conditional probabilities is given by

$$\begin{array}{c|ccc} & 0 & 1 & 2 \\ \hline 0 & p & q & 0 \\ 0^c & 0 & p & q \end{array}$$

e.g., $P(\text{Eve receives } 1 \mid \text{Alice sent message}) = p$. From this matrix and the fact that $P(A=0)=x$, one derives [1] that the capacity $C = \max_x \{ -px \log(px) - [qx + p(1-x)] \log [qx + p(1-x)] - q(1-x) \log [q(1-x)] + p \log(p) + (1-p) \log(1-p) \}$. C is maximized at one, when $p = 0$ or 1 , and C is minimized at .5 when $p = .5$.

4. Existing Anonymity Models

We do not deeply explore the existing models of anonymity for the sake of brevity and so that we may focus on explaining to our concepts of quasi-anonymity and covert channels. The existing models use different set-ups of mixes in their anonymity networks. Yet, when many of the existing mathematical analyses of mixes are examined, the themes are quite similar. In future work, we hope to apply our mathematical analyses to the entire literature of mixes and see how it holds up. We choose to extract the mathematical ‘nut’ of the existing models and compare them to ours. This is not to say that our approach is correct and the others are wrong. Rather, we are drawing a comparison between the different ways of modeling anonymity and the fact that we feel that some issues have been missed in the past models. We also briefly discuss the existing models to show that our thinking grew from the current literature and discussions in the field of anonymity. In future work, we will compare the different models and the different mix scenarios.

In [10], where discussing Crowds, the *degree of anonymity* is defined. This degree ranges from absolute

privacy to provably exposed. We are interested in their penultimate degree of anonymity defined as exposed. The degree of anonymity in the Crowds system is not mix based, but is determined by a probability p_f of a message being forwarded.

In [11] anonymity is quantified as $\log(N)$, where N is the number of senders. This is simply the size of the anonymity set in bits ([7] also discusses this). From this we see that as N increases so does the anonymity. However, it is too static a measurement to cope with dynamic observations by Eve.

In [12] Shannon's entropy is used as a measurement of anonymity where the concept of *effective size* S is discussed. This model has varying sender probabilities. When $S=0$, we are in a state of no anonymity. When S is maximized at $\log(N)$, Serjantov and Danezis argue that we are in a state of perfect anonymity. However, as our examples show, there are cases that fall outside of their model that have maximal entropy, but only quasi-anonymity.

We see a natural progression in [10], [11], and [12] from probability, to a normalized logarithmic count, to the entropy. Diaz *et al.* [13] took the next step in attempting to normalize the entropy and thus define the degree of anonymity as a number between 0 and 1. Unfortunately, this normalization is too gross a filter and removes the important factor of the bigger the anonymity set, the easier it is to hide. Instead, we crib from Shannon and consider the difference between the *a priori* entropy and the *a posteriori* conditional entropy, thus arriving at the mutual information. Of course, the scenario that we presented in the prior section is different than the threshold mixes discussed in most of the literature. We feel, though, that in light of the arguments presented in this paper, that the existing models should be reexamined in terms of a mutual information theory type approach. In our initial example (Fig. 1), simply maximizing the entropy of the senders does not give us anonymity. We can still pass information from Alice to Eve. This duality between Alice passing information to Eve, and the amount of quasi-anonymity is key to our thinking. We summarize below:

If the mutual information between Alice and Eve is non-zero, we only have quasi-anonymity.

In some cases we can further refine this by using the capacity (maximum mutual information). If we had perfect anonymity, Eve would never have any knowledge (deterministic or probabilistic) of Alice's behavior. In this case, Alice could not pass any information to Eve via the quasi-anonymous channel.

However, channel capacity alone is too broad a tool to use to measure all quasi-anonymous channels. For anonymity to be less than perfect, all that is required is for any small identifying piece of information to be leaked. In contrast, channel capacity is an asymptotic measurement; slight

temporal perturbations in system behavior have no effect upon capacity. A more general approach to the anonymity/covert channel correspondence considers the *mutual information* at each usage of the quasi-anonymous communication. This is arguably a better metric. For anonymity to be lost, one only has to discover a sender/receiver once. Therefore, an instantaneous metric such as mutual information (difference in entropy) may be better than an asymptotic approach. Of course, in the situation given above (discrete memoryless channel), there is no difference since the leakage of anonymity is there throughout the lifetime of the anonymity network. The examples that we have used discussed so far are based on a DMC model. It is possible for temporal variations in system behavior force to look at metrics using the more general form of mutual information that requires us to analyze the quasi-anonymous channel as a stochastic process. In particular, [5] gives an example of a covert channel with zero capacity, yet in this case Alice could pass any bit string to Eve in a noiseless manner. (Allow the time between transmission to increase exponentially, while at the same time only send one bit per transmission.) This uses Shannon's asymptotic definition of capacity of a noiseless channel as the upper limit of the ratio of the log of the number of symbols passed divided by the time.

We note an interesting recent paper by Danezis [14] where entropy is not discussed at all. Rather, he shows how the actual distribution may be analyzed if one has a long enough snapshot of the system. He proposes using hypothesis testing to determine the actual distribution of the senders. His approach is worthy of more analysis.


5. Quasi-anonymous Channel Examples

Even if the reader is not comfortable with using quasi-anonymous channels as a measure of anonymity, we still ask that the reader consider them as a threat in a quasi-anonymous network from a malicious user who may be in a trusted situation. The ability to leak information is not always accomplished by the easy method of "sneaker net." Or, it might not be users at all, but rather malicious software placed into a trusted high-assurance system that is leaking information out through a covert channel [4], [5]. Therefore, in the rest of this paper we concentrate on exploiting quasi-anonymity to leak information, or to allow an illicit communication between certain users.

Dual Enclave example revisited: We assumed that Eve only had knowledge of how many potential senders there were in Enclave1 and the outgoing message count, and that Alice and the C_i acted independently of each other. Let us remove the assumption that Alice has no knowledge of what the C_i are doing, and can act accordingly. This is realistic if Alice can tap what is going into the first mix firewall and react in time. If Alice has total knowledge of what the other C_i are doing, Alice can

force the parity¹ of the messages that Eve receives to be 0 or 1 noiselessly (Alice simply sends, or does not send a message to affect the parity). In this case Eve can receive 1 bit per t error-free, yet as long as there is at least one C_i that sends and at least one C_i that does not send, Eve will never know when, or if, Alice ever sent a message at all – implying that a covert channel may exist even in the presence of perfect anonymity. Note one could make this example, and the prior example, totally anonymous and wipe out the quasi-anonymous channel by padding [2], [15] what is coming out of the first mix firewall. Unfortunately, padding comes at the cost of system performance.

Exit only mix firewall: There is only one timed mix acting as a firewall now. This aside, the assumptions are as in the dual enclave scenario. In this example Eve knows how many messages each receiver (there are M of them) is getting (and of course the size of Enclave1). R_0 corresponds to a dummy receiver and is used to count messages not sent to any receiver. We see in this case that Eve's symbols are no longer $\{0, 1, \dots, N+1\}$, but rather they are all the ways to partition $N+1$ things into $M+1$ bins. Also keep in mind that the input alphabet is no longer $\{0, 0^c\}$. Rather, it



is $\{0, 1, \dots, M\}$. We see that the amount of information that can be sent over the quasi-anonymous channel is much greater than before. Also, if Alice is allowed to know of the C_i activity, capacity can be quite high indeed.

Threshold mix: As we noted earlier, mixes are often not used as firewalls. Rather they are used to simply hide the correspondence between sender and receiver, and Eve is allowed to know who is transmitting (that is, sender anonymity may not be a concern). As discussed, a threshold mix sits between the senders and the receivers. When a certain threshold K is met, the mix flushes the messages to the receivers. The anonymity in this scenario can be compromised by the $K-1$ attack [9], [11], where $K-1$ messages are sent by an attacker to see where a benign transmitter is sending messages. We see the lack of anonymity when using only one threshold mix. Similarly, there are many quasi-anonymous channels that arise. If Alice is a transmitter, she may certainly send information to Eve by sending or not sending a message. Eve can see that fact in the clear. But Alice can send more than this symbol. Alice can also send to different receivers. Eve is not sure which receiver Alice sent to, but as in the case of

the dual enclave, Alice's activity will influence what Eve sees coming out of the mix. If we had perfect sender/transmitter anonymity, there would be no such quasi-anonymous channels (and the $K-1$ attack would not work either).

Dual enclave threshold mix: Here Alice can influence the time (Eve's symbol) that the mix flushes. The mix will not flush itself until $K > 1$ messages have entered it. The time that the mix flushes of course depends on what the other members of the sending enclave do. Still assuming that a sender can send at most one message per t to the outgoing mix firewall, we have a quasi-anonymous *timing channel* from Alice to Eve.

If Alice is the only sender in the sending enclave, the first symbol she can send to Eve is Kt . Alice can do this by sending a message to the outgoing mix every t . When K messages have arrived the mix flushes. Now if Alice chooses not to send a message for one t , but sends for every other t , Eve will receive $Kt+1$; similarly, Alice can send Eve $Kt+2$, $Kt+3$, The capacities of such noiseless "timing channels" have been studied in [16]. If Alice only manipulates the mix to flush at Kt and $Kt+1$, this results in the lowest such capacity. As Alice ups the alphabet to $Kt+n$, the capacity increases. If Alice's symbol time n is unconstrained, the capacity is maximized. In general the capacity is bounded between $C_{T(Kt, Kt+1)}$ and $C_{T(Kt, 1)}$. Where $C_{T(a,b)}$ is given by the log of the positive root of $1 - (x^{-a} + x^{-b}) = 0$.

If Alice is not alone, and there are other senders sending messages to the outgoing mix, we then have a noisy timing channel. The capacity will, of course, be less than the noiseless case given above.

If we had perfect anonymity, and not the quasi-anonymity provided by the threshold mixes, we would have quasi-anonymous channels with zero capacity.

Dual enclave general Timing mix: Now we relax our initial assumption that the mix flushes every t . The mix may flush after a certain time interval $T > 1$ (independent of what is in the mix buffer). Of course, we are still in the situation of a noisy timing channel, but the effects of Alice are moderated by the other senders and by the forced latency in sending messages.

Pool mix: Note that in any of these situations we could also use a pool mix which would hold back a certain number of messages each time the mix flushes. This causes more noise to be introduced into the quasi-anonymous channel.

Pump mix: In [17], [18] the idea of a Pump was put forth as a way to greatly lessen the threat of certain covert channels. We put forth the idea of a Pump mix to lessen the capacity of quasi-anonymous channels and to increase the anonymity of a system. The Pump mix would record

¹ We thank Christopher Lynch for this observation.

the history of senders input record to the mix. If one particular sender is “hogging” the mix, the Pump mix would delay messages from that sender. This would also have the effect of enforcing a fairness policy across the anonymity network. The delay would be implemented as in standard Pump algorithm by using a random variable and a moving average of the past history. Pump mixes are related to Stop-and-Go mixes [19], see also [20] and the Cottrell mix [9]. However, no messages would be dropped in a Pump mix; they would be stored in a stable buffer. Of course this assumes an arbitrarily large buffer.

6. Conclusions

We have linked the lack of complete anonymity to that of covert communication. We have demonstrated that covert channels can both leak information and be used as a metric for quasi-anonymity. Additionally, we have examined various mix architectures and have identified various covert channels.

Concepts such as covert channels have been used to show weaknesses in high-assurance systems for many years. The roots of covert channels are in information theory. There has been research extending this notion to specific formal models such as non-interference [21]. In fact, there is recent work extending probabilistic non-interference [22]. However, every model still has its roots in Shannon [6]. Anonymity is a rather new field. It will take time for the definitions, concepts, and devices, as set forth in [7], to get sorted out. We hope this paper furthers that aim.

7. Acknowledgements

We thank R. Heilizer, C. Lynch, and the anonymous reviewers. Research funded by ONR.

References

- [1] I. S. Moskowitz, R.E. Newman, D.P. Crepeau & A.R. Miller, Covert channels and anonymizing networks, *Proc. WPES 2003*, Washington, DC, pp. 79-88, 2003.
- [2] B.R. Venkatraman & R.E. Newman-Wolfe, Capacity estimation and auditability of network covert channels, *Proc. Symp. on Security & Privacy*, CA, 186-198, 1995.
- [3] R.E. Newman, I.S. Moskowitz, P. Syverson & A. Serjantov, Metrics for traffic analysis prevention, *Proc. PET 2003*, Dresden, March 2003.
- [4] B.W. Lampson, A note of the confinement problem, *Communications of the ACM*, 16(10), Oct. 1973.
- [5] I.S. Moskowitz & M.H. Kang, Covert channels --- here to stay?, *Proc. COMPASS '94*, MD, 1994, 235-243.
- [6] C.E. Shannon, The mathematical theory of communication, *Bell Systems Tech. J.*, 30:50-63, 1948.
- [7] A. Pfitzmann & M. Köhntopp, Anonymity, unobservability and pseudonymity --- a proposal for terminology. *Proc. Designing Privacy Enhancing Technologies*, LNCS 2009, pp. 10-29, Springer 2001. (extended through 27 May 2003)
- [8] D. Chaum, Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2), pp. 84-88, 1981.
- [9] A. Serjantov, R. Dingledine, & P. Syverson, From a trickle to a flood, *Proc. Information Hiding '02*, LNCS 2578, pp. 36-52, Springer 2003.
- [10] M.K. Reiter & A.V. Rubin, Crowds: anonymity for web transactions, *ACM Transactions on Information and Systems Security*, 1(1), pp. 66-92, Nov. 1998.
- [11] O. Berthold, A. Pfitzmann, & R. Standtke, The disadvantages of free MIX routes and how to overcome them, *Proc. Designing Privacy Enhancing Technologies*, LNCS 2009, pp. 30-45, Springer 2001.
- [12] A. Serjantov & G. Danezis, Towards an information theoretic metric for anonymity, *Proc. PET 2002*, LNCS 2482, pp. 41-53, Springer 2002.
- [13] C. Diaz, S. Seys, J. Classens, & B. Preneel, Towards measuring anonymity, *Proc. PET 2002*, LNCS 2482, pp. 54-68, Springer 2002.
- [14] G. Danezis, Statistical disclosure attacks-traffic confirmation in open networks, pp. 421-426, *Proc. SEC2003*.
- [15] O. Berthold, H. Federrath, & S. Köpsell, Web MIXes: a system for anonymous and unobservable internet access, *Designing Privacy Enhancing Technologies*, LNCS 2009, pp. 115-129, Springer 2001.
- [16] I. S. Moskowitz & A.R. Miller, Simple timing channels, *Proc. IEEE Symp. On Security and Privacy*, Oakland, Ca. pp. 56-64, May 1994.
- [17] M.H. Kang & I.S. Moskowitz, A Pump for rapid, reliable, secure communication, *Proc. ACM Conf. On Computer & Communication Security*, Fairfax, VA, pp. 119-129, 1993.
- [18] M.H. Kang, I.S. Moskowitz, & D.C. Lee, A network version of the Pump, *Proc. IEEE Symp. on Security & Privacy*, CA, pp. 144-154, May 1995.
- [19] D. Kesdogan, J. Egner, & R. Büschkes, Stop-and-Go-MIXes providing probabilistic anonymity in an open system, *Proc. Information Hiding '98*, LNCS 1525, pp. 83-98, Springer 1998.
- [20] J.-F. Raymond, Traffic analysis: protocols, attacks, design issues, and open problems, *Proc. Designing Privacy Enhancing Technologies*, pp. 10-29, 2000.
- [21] J.A. Goguen & J. Meseguer, Unwinding and inference control, *Proc. IEEE Symp. On Research in Security and Privacy*, Oakland, CA, pp. 75-86, 1987.
- [22] M. Backes & B. Pfitzmann, Computational probabilistic non-interference, *Proc. ESORICS 2002*, LNCS 2502, pp. 1-15, Springer, 2002.